Recoding Color Transfer as a Color Homography

Han Gong[1]
http://www2.cmp.uea.ac.uk/~ybb15eau/

Graham D. Finlayson[1]
g.finlayson@uea.ac.uk

Robert B. Fisher[2]
http://homepages.inf.ed.ac.uk/rbf/

[1] School of Computing Sciences
University of East Anglia
Norwich, UK

[2] School of Informatics
University of Edinburgh
Edinburgh, UK

Abstract

Color transfer is an image editing process that adjusts the colors of a picture to match a target picture's color theme. A natural color transfer not only matches the color styles but also prevents after-transfer artifacts due to image compression, noise, and gradient smoothness change. The recently discovered color homography theorem proves that colors across a change in photometric viewing condition are related by a homography. In this paper, we propose a color-homography-based color transfer decomposition which encodes color transfer as a combination of chromaticity shift and shading adjustment. A powerful form of shading adjustment is shown to be a global shading curve by which the same shading homography can be applied elsewhere. Our experiments show that the proposed color transfer decomposition provides a very close approximation to many popular color transfer methods. The advantage of our approach is that the learned color transfer can be applied to many other images (*e.g.* other frames in a video), instead of a frame-to-frame basis. We demonstrate two applications for color transfer enhancement and video color grading re-application. This simple model of color transfer is also important for future color transfer algorithm design.

1 Introduction

Adjusting the color style of pictures/frames is one of the most common tasks in professional photo editing as well as video post-production. Artists would often choose a desired target picture and manipulate the other pictures to match their target color style. This process is called color transfer. An example of color transfer between a source image and a target image is shown in Figure 1. Typically, this color tuning process requires artists to delicately adjust for multiple properties such as exposure, brightness, white-point, and color mapping. These adjustments are also interdependent, *i.e.* when aligning an individual property this may cause the others to become misaligned. For rendered images, some artifacts (*e.g.* JPEG block edges) may appear after color adjustment. It is therefore desirable to automate this time-consuming task.

Example-based color transfer was first introduced by Reinhard *et al*. [18]. Since then, much further research [13, 15, 16, 17] has been carried out. A recent discovery of the color

homography theorem reveals that colors across a change in viewing condition (illuminant, shading or camera) are related by a homography [6, 7]. In this paper, we propose a general model, based on the color homography theorem, to approximate different color transfer results. In our model, we decompose any color transfer into a chromaticity mapping component and a shading adjustment component. We also show that the shading adjustment can be reformulated by a global shading curve through which the shading homography can be applied elsewhere. Our experiments show that our model produces very close approximations to the original color transfer results. We believe that our color transfer model is useful and fundamental for developing simple and efficient color transfer algorithms. That is a trained model can be applied to every frame in a video fragment rather than needing frame-by-frame adjustment. This decomposition also enables users to amend the imperfections of a color transfer result or simply extract the desired effect.

Our paper is organized as follows. We review the popular color transfer methods and the color homography theorem in §2. Our color transfer decomposition is described in §3. We show our evaluation and applications in §4. Finally, we conclude in §6.

2 Background

2.1 Color transfer

Example-based color transfer was first introduced by Reinhard *et al*. [18]. Their method assumes that the color distribution in $l\alpha\beta$ color space is a normal distribution. They map a source image to its target so that their color distributions have the same mean and variance in $l\alpha\beta$ color space. Pitie *et al*. [15] proposed an iterative color transfer method that rotates and shifts the color distribution in 3D until the distributions of the two images are aligned. The rotation matrix is random over all possible angular combinations. This method was later improved by adding a gradient preservation constraint to reduce after-transfer artifacts [16]. Pouli and Reinhard [17] adopted a progressive histogram matching in L*a*b* color space. In their color transfer method, users can specify the level of color transfer (i.e. partial color transfer) between two images. Their algorithm also addresses the difference in dynamic ranges of between two images. Nguyen *et al*. [13] proposed an illuminant-aware and gamut-based color transfer. A white-balancing step is first performed for both images to remove color casts caused by different illuminations. A luminance matching is later performed by histogram matching along the "gray" axis of RGB. They finally adopt a 3D convex hull mapping, which contains scale and rotation operations, to ensure that the color-transferred RGBs are still in the space of the target RGBs. There are some other approaches (*e.g.* [1, 3, 10, 19, 21]) that solve for several local color transfers rather than a single global color transfer. In this paper, we focus on global color transfer.

Pitie *et al*. [14] proposed a color transfer approximation by a 3D affine mapping. The linear transform minimizes the amount of changes in color and preserves the monotonicity of intensity changes. However, it is based on the assumption that the color distributions of the two images are both normal distributions. It also does not well approximate the shading change of a color transfer.

2.2 Color homography

The color homography theorem [6, 7] shows that chromaticities across a change in capture conditions (light color, shading and imaging device) are a homography apart. Let us map an RGB $\underline{\rho}$ to a corresponding RGI (red-green-intensity) $\underline{c}$ using a 3×3 full-rank matrix C:

$$\underline{\rho}^{\mathsf{T}} C = \underline{c}^{\mathsf{T}}$$

$$
\begin{bmatrix} R \\ G \\ B \end{bmatrix}^{\mathsf{T}}
\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}
=
\begin{bmatrix} R \\ G \\ R+G+B \end{bmatrix}^{\mathsf{T}}
\tag{1}
$$

The r and g chromaticity coordinates are written as $r = R/(R+G+B)$, $g = G/(R+G+B)$. We interpret the right-hand-side of Equation 1 as a homogeneous coordinate and we have $\underline{c} \propto \begin{bmatrix} r & g & 1 \end{bmatrix}^{\mathsf{T}}$. When the shading is fixed, it is well-known that across a change in illumination or a change in device, the corresponding RGBs are related by a 3×3 linear transform M that $\underline{\rho}^{\mathsf{T}} M = \underline{\rho}'^{\mathsf{T}}$ where $\underline{\rho}'$ is the corresponding RGBs under a second light or captured by a different camera [1, 2]. Clearly, $H = C^{-1}MC$ maps colors in RGI form between illuminants. Due to different shading, the RGI triple under a second light is written as $\underline{c}'^{\mathsf{T}} = \alpha \underline{c}^{\mathsf{T}} H$, where α denotes the unknown scaling. Without loss of generality let us interpret $\underline{c}$ as a homogeneous coordinate i.e. assume its third component is 1. Then, $[r' \ g']^{\mathsf{T}} = H([r \ g]^{\mathsf{T}})$ (rg chromaticity coordinates are a homography $H()$ apart).

3 Recoding color transfer as a color homography

Color transfer can often be interpreted as re-rendering an image with respect to real physical scene changes (e.g. from summer to autumn) and/or illumination. Recent work [14] approximates the effect of global color transfer by a 3D affine mapping. We propose that, in general, we can better approximate most global color transfer algorithms as a color homography transfer. The color homography theorem shows that the same scene under an illuminant (or camera) change will result in two images a homography apart. We propose that a global color transfer can be decomposed into a linear chromaticity mapping and a shading adjustment. This concise form enables us to efficiently replicate the originally slow color transfer and re-apply it to many other images (*e.g.* the frames of a video fragment).

Throughout the paper we denote the source image by I_s and the original color transfer result by I_t. Given I_s and I_t, the aim of color transfer decomposition is to find a general model that reproduces the color theme change from I_s to I_t. Figure 1 shows our two-step color transfer decomposition: 1) Chromaticity mapping estimation (*simple* homography). The source image is chromaticity transferred by applying a color homography transform estimated from the corresponding rg chromaticities of the source image and the original color transfer result. 2) Shading adjustment estimation (*shading* homography). A further shading adjustment is estimated by finding a least-squares solution that aligns the shadings of the original color transfer output and the chromaticity transferred image. These procedures are explained in detail in the following.
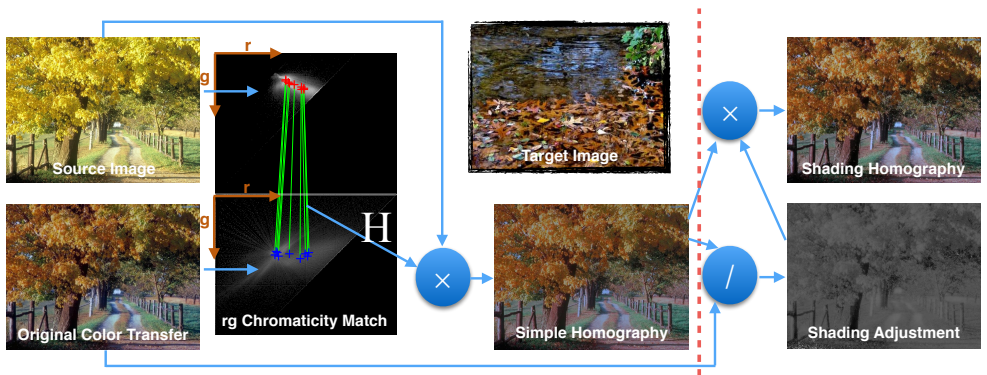
Figure 1: Pipeline of color-homography-based color transfer decomposition. The red dashed line divides the pipeline into two steps: 1) *Simple* homography. The rg chromaticities of the source image and the original color transfer image (by [16]) are matched according to their chromaticity locations (*e.g.* the green lines), from which we estimate a color homography matrix H and use H to transfer the source image. 2) *Shading* homography. The shadings are aligned between the *simple* homography result and the original color transfer result by a least-squares method. The per-pixel product of the *simple* homography result and the shading adjustment gives a close color transfer approximation.

3.1 Color Homography Color Transfer Model

We start with the outputs of the prior-art algorithms. Assuming we relate I_s to I_t with a pixel-wise correspondence, we represent the RGBs of I_s and I_t as two $n \times 3$ matrices A and B respectively where n is the number of pixels. These $n \times 3$ matrices can be reconstituted into the original image grids. The chromaticity mapping is modeled as a 3×3 linear transform but because of the relative positions of light and surfaces there might also be per-pixel shading perturbations. Assuming the Lambertian image formation an accurate physical model,

$$DAH \approx B \qquad (2)$$

where D is an $n \times n$ diagonal matrix of shading factors and H is a 3×3 chromaticity mapping matrix. A color transfer can be decomposed into a diagonal shading matrix D and a homography matrix H. The homography matrix H is a global chromaticity mapping. The matrix D can be seen as a change of surface reflectance or position of illuminant.

According to the color homography model, we define two color transfer decomposition models. In *simple* homography transfer, the output image is a homography from the input which only contains a chromaticity mapping. It preserves the shading of the source image and does not include the shading of the original color transfer result. In *shading* homography transfer, the output also incorporates the best shading factors which restore the shading of the original color-transfer output. By solving for the homography H, the *simple* and *shading* homography transfers are defined as:

$$B_{simple} = AH \approx B \qquad (3)$$

$$B_{shading} = DAH = DB_{simple} \approx B \qquad (4)$$

An example of the two color transfer models are shown in Figure 1. When recoding a color transfer, these two recoding models provide an additional flexibility as some users may only attempt to extract the chromaticity mapping.

3.2 Chromaticity mapping estimation

In color correction, Equation 2 is solved by using Alternating Least-Squares (ALS) [6, 7, 8] illustrated in Algorithm 1.

---

1   $i = 0$, $\min_{D^0} \left\| D^0 A - B \right\|_F$, $A^0 = D^0 A$;
2   **repeat**
3     $i = i + 1$;
4     $\min_{H^i} \left\| A^{i-1} H^i - B \right\|_F$;
5     $\min_{D^i} \left\| D^i A^{i-1} H^i - B \right\|_F$;
6     $A^i = D^i A^{i-1} H^i$;
7   **until** $\left\| A^i - A^{i-1} \right\|_F < \varepsilon$;

Algorithm 1: Homography from alternating least-squares

The effect of the individual H^i and D^i can be merged into a single matrix $D = \prod_i D^i$ and $H = \prod_i H_i$ (where the product is taken by post-multiplying matrices). $\|.\|_F$ denotes the Frobenius norm and H^i and D^i are found using the closed form Moore-Penrose inverse. In color transfer, we choose to minimize the rg chromaticity (*i.e.* normalized RG) difference between two images because any non-zero RGB can be mapped to the range of $[0, 1]$. To achieve this, we modify Equation 2 as

$$Dh(AC)H_{rg} \approx h(BC) \tag{5}$$

where C is the 3×3 RGB-to-RGI conversion matrix defined in Equation 1, h is a function that converts each RGI intensity (matrix row) to their homogeneous coordinates (by dividing RGI by I which makes all elements of the 3^{rd} column 1), H_{rg} is the homography matrix that minimizes rg chromaticity difference. The homography matrix H for color transfer is related with H_{rg} by $H = CH_{rg}C^{-1}$. As the under-saturated pixels with zero RGBs contain no color information, they are excluded from the computation.

To reduce the computational cost, it is possible to estimate H with down-sampled images. We find that image down-sampling barely affects our chromaticity mapping quality. An example is shown in Figure 2. Depending on the content of image, the minimum down-sampling factor for estimating H may vary.

3.3 Shading adjustment estimation

A chromaticity-transferred result may still not be close to the actual color transfer result because color transfer methods also adjust contrast and intensity mapping. In our approximation pipeline, the shading adjustment matrix D can be directly obtained from the ALS procedure. When the chromaticity mapping matrix H is estimated from down-sampled images, the estimated D from ALS is not in full-resolution. In this case, according to Equation 4, D can be solved for by a least-squares solution $\min_D \left\| DB_{simple} - B \right\|_F$. We introduce the additional shading reproduction step as follows.

Figure 2: Down-sampled images for chromaticity mapping estimation (shading adjustment is still estimated using full-resolution images). The sizes of I_s and I_t are reduced by the corresponding factors. Image down-sampling barely affects the color transfer approximation quality such that even two 3×4 thumbnails are sufficient for getting the approximation result shown in the right-most figure (2^{-8} down-sampling).

Shading adjustment reproduction – mapped shading homography
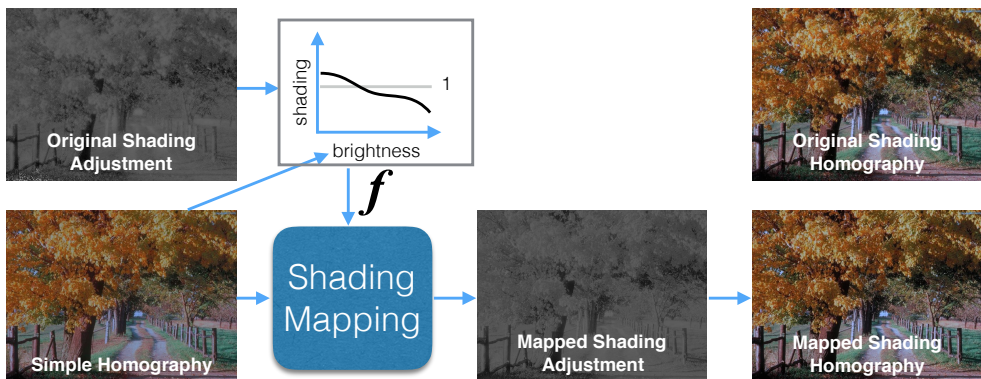

Figure 3: Shading reproduction. A function f of brightness-to-shading mapping is fitted to the brightness of the simple homography result and the original shading adjustment. The function f is used to reproduce a mapped shading adjustment from which a mapped shading homography result is generated.

A universal color transfer decomposition should be compatible with any input image in a similar color theme. Although the chromaticity mapping matrix H is reusable for adjusting other images, the per-pixel shading adjustment D derived from the ALS procedure only works for the source image. To resolve this issue, we propose a brightness-to-shading mapping to reproduce the shading adjustment for any input image. The mapping is estimated by fitting a smooth curve to the per-pixel brightness and shading data. A direct fitting for all data points is computationally costly. Instead, we uniformly divide the brightness range into 50 slots and compute the center point (average brightness and shading) for each slot. The center point summarizes the point distribution of a range to reduce the amount of data for fitting. The smooth piece-wise curve f is modeled as a Piece-wise Cubic Hermite Interpolating Polynomial (PCHIP) [9] according to the 50 summarized data sites.

Directly applying the mapped shading adjustment may lead to sharp gradient artifacts because the shading variations for some areas may not follow the global trend. When the shading is not smooth, the overall magnitude of image edges is expected to be large. Inspired by the Laplacian smoothness constraint adopted in [5], we solve this by minimizing the

overall magnitude of Laplacian image edges of the shading image as shown in Equation 6:

$$\min_D \left\| D - D_{\text{mapped}} \right\|_F + \lambda \left\| I_D * K \right\|_F \tag{6}$$

where the first term ensures the similarity between the optimum shading D and the f-mapped per-pixel shading D_{mapped}, the second term enforces the smoothness constraint, λ is a smoothness weight with a small value, I_D is the 2D shading image (reshaped from the vector of the diagonal elements of D) which is then convolved by a 3×3 Laplacian kernel K. See our supplementary material for the details about how to solve this minimization and determine λ adaptively. As is shown in Figure 3, the shading adjustment can be reproduced according to f and the brightness of the *simple* homography transfer result. The result generated from mapped shading is visually close to the original *shading* homography result.

4 Results

We first show some visual results of color transfer approximations of [13, 16, 17, 18] in Figure 4. Global 3D affine mapping [14] does not well reproduce the shading adjustments of color transfer. Our homography-based method offers a closer color transfer approximation. In Table 1, we also quantitatively evaluate the approximation accuracy of 3 candidates by a PSNR (Peak Signal-to-Noise Ratio) measurement. Acceptable values for wireless image transmission quality loss are considered to be over 20 dB (the higher the better) [20]. The test is based on 7 classic color transfer image pairs and 4 color transfer methods. The original shading homography produces the best result overall. Mapped shading homography also generally produces higher PSNR scores compared with 3D affine mapping, *esp.* for [16, 17, 18].

Table 1: PSNR measurement between the original color transfer result and its approximation (see our supplementary material for the complete table and their visual results).

	Nguyen [13]	Pitie [16]	Pouli [17]	Reinhard [18]
3D affine [14]	26.85	26.04	26.92	28.49
Shading homography	31.51	31.06	36.55	35.48
Mapped shading homography	27.77	28.16	31.70	31.18

5 Applications

In this section, we show that our color transfer decomposition can be applied to color transfer enhancement and video color grading re-application.

Color transfer with reduced artifacts

In Figure 5, original shading homography approximation retains the artifacts of noise and over-saturation of the original color transfer result. These artifacts are fixed by mapped *shading* homography.
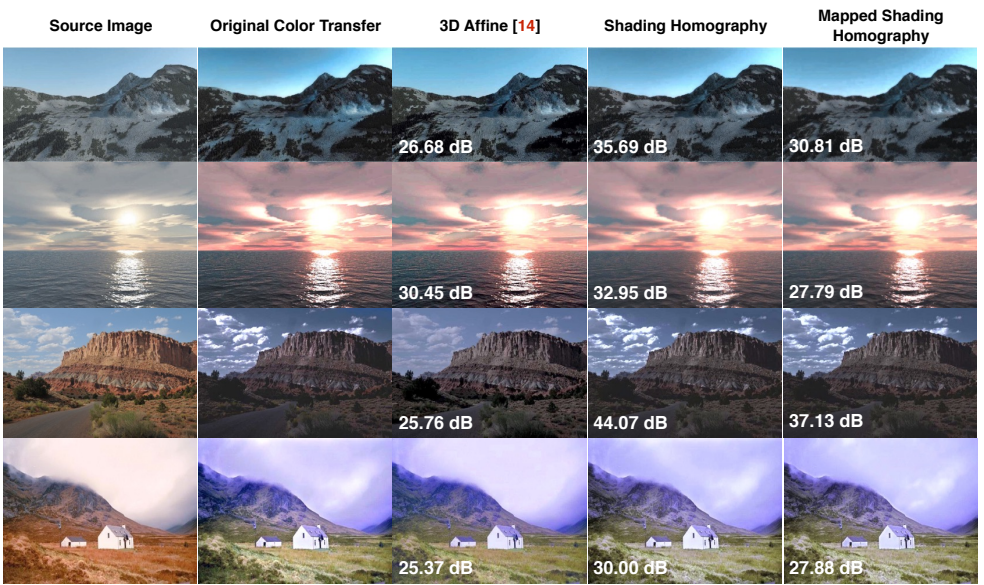
Figure 4: Visual result of color transfer approximations (in the order of [16], [18], [17], [13]). The corresponding PSNR error is shown on each approximation result. Our homography-based methods produce closer approximations to the original color transfer results.
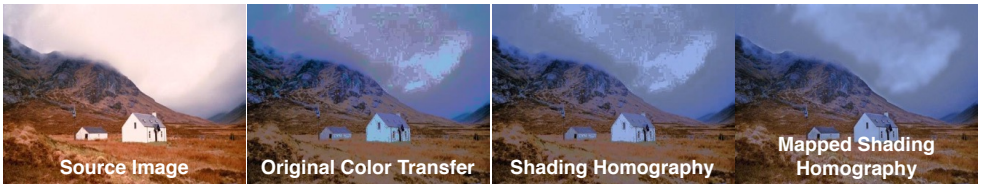

Figure 5: Imperfection fixing. The imperfections in a color transfer [17] and its shading homography approximation are fixed by mapped shading adjustment.

Video color grading re-application

Color transfer methods [13, 16, 17, 18] cannot be directly applied to video color grading as per-frame color matching leads to temporal incoherence [2, 4]. The color homography model is a concise representation of the original complex video color grading adjustments. Compared with a prior art [2], our model generates stable results in one-go without the excessive steps for removing artifacts such as flickering and bleeding. As shown in Figure 6, given two sample images profiling the desired color grading adjustment, the complex steps of video color grading can be extracted as a mapped *shading* homography transfer. The extracted color grading effect can also be re-applied to a different video sequence of a similar color theme (see our supplementary video for more examples).

Figure 6: Video color grading re-application (original video from Cry ©Jeffro). The color grading profile is extracted from two pairs of image samples. Grading profile 1 is applied to the same scene. Grading profile 2 is applied to a scene different from its image samples.

6 Conclusion

Based on the theorem of color homography, we have shown that a global color transfer can be approximated by a combination of chromaticity mapping and shading adjustment. Our experiment shows that the proposed color transfer decomposition approximates very well many popular color transfer methods. We also demonstrate two applications for fixing the imperfections in a color transfer result and video color grading re-application. We believe that this verified model of color transfer is also important for developing simple and efficient color transfer algorithms.

Acknowledgment

References

[1] Xiaobo An and Fabio Pellacini. User-controllable color transfer. In *Computer Graphics Forum*, volume 29, pages 263–271. Wiley Online Library, 2010.

[2] Nicolas Bonneel, Kalyan Sunkavalli, Sylvain Paris, and Hanspeter Pfister. Example-based video color grading. *ACM Transactions on Graphics*, 32(4):39–1, 2013.

[3] Huiwen Chang, Ohad Fried, Yiming Liu, Stephen DiVerdi, and Adam Finkelstein. Palette-based photo recoloring. *Transactions on Graphics*, 34(4):139, 2015.

[4] Zeev Farbman and Dani Lischinski. Tonal stabilization of video. *ACM Transactions on Graphics*, 30(4):89, 2011.

[5] Zeev Farbman, Raanan Fattal, Dani Lischinski, and Richard Szeliski. Edge-preserving decompositions for multi-scale tone and detail manipulation. In *ACM Transactions on Graphics*, volume 27, page 67. ACM, 2008.

[6] Graham D. Finalyson, Han Gong, and Robert B. Fisher. Color homography color correction. In *Color Imaging Conference*. Society for Imaging Science and Technology, 2016.

[7] Graham D. Finalyson, Han Gong, and Robert B. Fisher. Color homography. In *Progress in Colour Studies*. John Benjamins Publishing Company, 2016.

[8] Graham D. Finlayson, Maryam Mohammadzadeh Darrodi, and Michal Mackiewicz. The alternating least squares technique for non-uniform intensity color correction. *Color Research & Application*, 40(3):232–242, 2015.

[9] Frederick N Fritsch and Ralph E Carlson. Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246, 1980.

[10] Sefy Kagarlitsky, Yael Moses, and Yacov Hel-Or. Piecewise-consistent color mappings of images acquired under various conditions. In *International Conference on Computer Vision*, pages 2311–2318. IEEE, 2009.

[11] L.T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, 3:1673–1683, 1986.

[12] D.H. Marimont and B.A. Wandell. Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A*, 9(11):1905–1913, 1992.

[13] R. M. H. Nguyen, S. J. Kim, and M. S. Brown. Illuminant aware gamut-based color transfer. *Computer Graphics Forum*, 33(7):319–328, October 2014.

[14] F Pitié and A Kokaram. The linear monge-kantorovitch linear colour mapping for example-based colour transfer. In *European Conference on Visual Media Production*, pages 1–9. IET, 2007.

[15] F. Pitie, A.C. Kokaram, and R. Dahyot. N-dimensional probability density function transfer and its application to color transfer. In *International Conference on Computer Vision*, volume 2, pages 1434–1439, 2005.

[16] François Pitié, Anil C. Kokaram, and Rozenn Dahyot. Automated colour grading using colour distribution transfer. *Computer Vision And Image Understanding*, 107(1-2):123–137, July 2007.

[17] Tania Pouli and Erik Reinhard. Progressive histogram reshaping for creative color transfer and tone reproduction. In *International Symposium on Non-Photorealistic Animation and Rendering*, pages 81–90, New York, NY, USA, 2010. ACM.

[18] Erik Reinhard, Michael Ashikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer Graphics And Applications*, 21(5):34–41, September 2001.

[19] Yu-Wing Tai, Jiaya Jia, and Chi-Keung Tang. Local color transfer via probabilistic segmentation by expectation-maximization. In *Conference on Computer Vision and Pattern Recognition*, pages 747–754. IEEE, 2005.

[20] Nikolaos Thomos, Nikolaos V Boulgouris, and Michael G Strintzis. Optimized transmission of jpeg2000 streams over wireless channels. *Transactions on Image Processing*, 15(1):54–67, 2006.

[21] Fuzhang Wu, Weiming Dong, Yan Kong, Xing Mei, Jean-Claude Paul, and Xiaopeng Zhang. Content-based colour transfer. In *Computer Graphics Forum*, volume 32, pages 190–203. Wiley Online Library, 2013.